determining that an image from a second viewpoint is to be rendered, the second viewpoint associated with a second viewing frustum;

determining that a first octant of the plurality of octants intersects the second viewing frustum;

receiving a plurality of points included in the first octant from a cache memory of a central processing unit coupled to the GPU;

replacing a point at a first position in the first frame buffer object with a first point of the plurality of points that represents the same pixel and that is closer to a position of the virtual camera;

determining that a gap exists in the first frame buffer object at a second position;

filling the second position in the first frame buffer object using a second point at a corresponding position in a second frame buffer object storing a plurality of points from the first synthetic point cloud;

determining that a first gap exists in a third frame buffer object storing a plurality of points from a third synthetic point cloud at the first position, and that a second gap exists in the third frame buffer object at a second position;

filling the first gap using the first point;

filling the second gap using the second point;

rendering a second image corresponding to the second viewing frustum based on points stored in the third frame buffer object, including the first point and the second point.

14. The method of claim 13, further comprising:

determining that a second octant of the plurality of octants intersects the second viewing frustum;

selecting the first octant based on a distance to the first octant;

requesting points in the first octant from the cache memory; and

receiving the plurality of points responsive to the request.

15. The method of claim 13, wherein the second point represents a first cell in the first synthetic point cloud, and a color of the second point is an average calculated based on the color of each of a plurality of points from the point cloud data that are at least partially within the portion of the scene represented by the first cell.

16. The method of claim 15, wherein the color of the second point is based on a simple average of the plurality of points from the point cloud data.

17. The method of claim 15, wherein a position of the second point is a midpoint of the first cell.

18. The method of claim 13, further comprising storing the third synthetic point cloud representing the scene, the third synthetic cloud subdivided into the plurality of octants, each of the plurality of octants subdivided into eight sub-octants that are each organized as a grid of third voxels, wherein the third voxels each represent a smaller portion of the scene than the second voxels.

19. The method of claim 13, wherein the first synthetic point cloud is subdivided into n voxels along a first axis of the scene, where n is less than or equal to 256, and the first octant is subdivided into n voxels along the first axis.

20. The method of claim 19, wherein a position of each point of the first synthetic point cloud within the first synthetic point cloud is represented by no more than three bytes.

21. The method of claim 13, wherein a position of each point of the second synthetic point cloud within an octant of the plurality of octants is represented using no more than three bytes.

* * * * *